RFC1918 Caching Security Issues

RFC1918 Caching Security Issues - Robert Hansen, sectheory.com.

- Published: date not stated
- Original: <http://www.sectheory.com/rfc1918-security-issues.htm>
- Preserved from: http://www.sectheory.com/rfc1918-security-issues.htm (manual-import) on 2026-09-10
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

RFC1918 Caching Security Issues

By Robert Hansen Date: 08/06/2009

Preface: Intranets are intended to be secured from the outside by way of firewalls and other networking devices. Unfortunately, there has been a move towards non publicly-routable address space as a method of protection, rather than other methods of protecting private IP space. This paper will outline a number of flaws that can be exploited by an adversary because of the use of well known non publicly-routable IP address spaces.

Overview: One of the principle technologies employed by enterprises is the concept of non publicly-routable IP address space (otherwise known as RFC1918). RFC1918 as defined explains that one of the principle reasons people use it is to avoid the future IP exhaustion that IPv6 is intended to obviate. Unfortunately, it accomplishes this task by using the same set of IP spaces for everyone who uses this tactic.

```
10.0.0.0        -  10.255.255.255  (10/8 prefix)
172.16.0.0      -  172.31.255.255  (172.16/12 prefix)
192.168.0.0     -  192.168.255.255 (192.168/16 prefix)
```

The bulk of intranet IP space falls in 10.* and 192.168*., and further, the bulk falls in 10.0.0.*, 10.1.1*., 10.10.10.*, 192.168.0.* and 192.168.1.*. Narrowing the most likely subnets down to 1280 addresses (256 addresses * 5 subnets). This tends to lead to collision of IP space where two separate networks will look virtually identical from an IP perspective. There are many technologies that use

private IP addresses as a method of securing themselves. Likewise the browsers have implemented the same origin policy which prohibits a server on the Internet from reading content on another server (this includes internal address space).

Because of caching issues within the browser, and other technologies that may use the IP address as the single factor of security, it becomes possible to create situations where the collisions can be used to an attacker's advantage, and even allow them to compromise internal networks.

The Attacks: There are a number of potential attacks that are possible, and many of them reside around trust relationships people have with third parties. One such instance is a VPN (virtual private network) connection between a good entity and one that intends to compromise the victim's network.

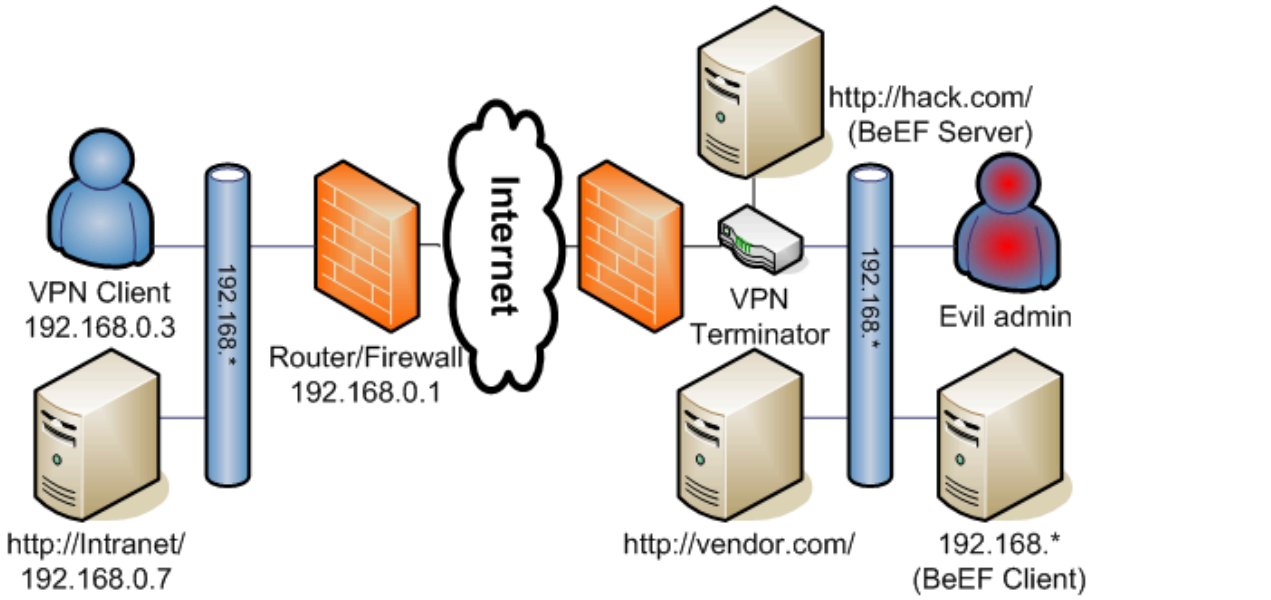

Step 1) Attacker does DNS recon and finds intranet server is called http://Intranet/
Step 2) User visits http://Intranet/ (DNS pinning occurs in the user's browser)
Step 3) User instantiates client VPN to Vendor's network
Step 4) Evil admin pushes all RFC1918 routes to connect to his own network
Step 5) User connects to http://vendor.com/ and renders an iframe to http://Intranet/
Step 6) User connects to evil network instead of the real Intranet page (also giving attacker access to the user's cookies)
Step 7) Evil admin delivers a BeEF client payload and requests that the user reconnect to the same page in several seconds.
Step 8) Evil admin turns off the VPN client forcing the routes to change back to their original settings (depending on the client).
Step 9) User connects to http://Intranet/ within his own network, but now under the control of the BeEF client payload.
Step 10) The BeEF client payload request C&C instructions from http://hack.com
Step 11) Upon successful connection of the BeEF client to the BeEF server the Evil Admin allows the user to reconnect to the VPN, but pushes only necessary routes.

Fig 1. Click to enlarge

The first attack, as seen in *Fig 1* is where a user is using a client VPN and connecting into a hostile network. The user's browser is thwarted into routing then visiting and caching many pages that would normally be reserved for internal addresses within their own network. Because of caching issues within browsers, there is no need to break the same origin policy - only to wait. Once the VPN connection is destroyed (assuming the routes are then broken) the user's browser then connects to their real RFC1918 addresses, which are now under the control of an attacker by way of a JavaScript back door (Eg: BeEF). This sort of exploit would be most often seen between two competitive companies that share information only occasionally, or between two companies in a partner/vendor relationship.

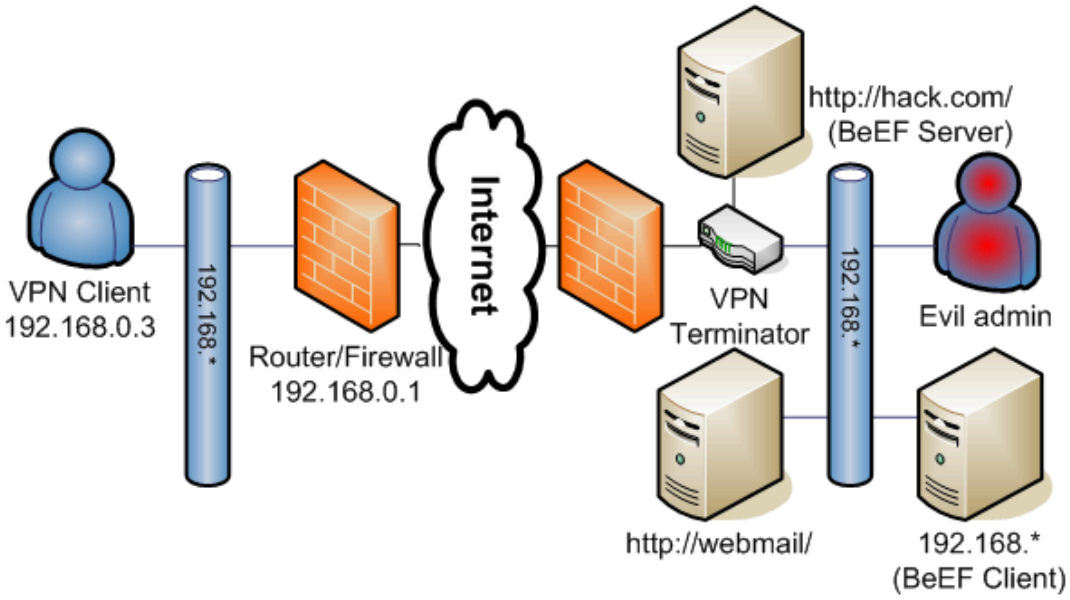

Step 1) User instantiates client VPN to work network
Step 2) Evil admin pushes all RFC1918 routes to connect to his own network
Step 3) User connects to http://webmail/ and renders an iframe to http://192.168.0.0/ 0.1 and 1.1
Step 4) User connects to evil network instead of the real router page (also giving attacker access to the user's cookies)
Step 5) Evil admin delivers a BeEF client payload and requests that the user reconnect to the same page in several seconds.
Step 6) Evil admin turns off the VPN client forcing the routes to change back to their original settings (depending on the client).
Step 7) User connects to internal addresses within his own network, but now under the control of the BeEF client payload.
Step 8) The BeEF client payload request C&C instructions from http://hack.com
Step 9) Upon successful connection of the BeEF client to the BeEF server the Evil Admin allows the user to reconnect to the VPN, but pushes only necessary routes.

Fig 2. Click to enlarge

The second attack as seen in *Fig 2*, similar to the first, requires instead of it being two office networks, it is a user who is using a client VPN from a home office. Like the first example, caching within the browser allows an evil administrator to persist their JavaScript backdoor beyond the lifetime of the VPN connection. The administrator in this way could compromise the user's home network. This may effect administrators who want to compromise executive management's home networks, without leaving as large a trail as traditional malware.

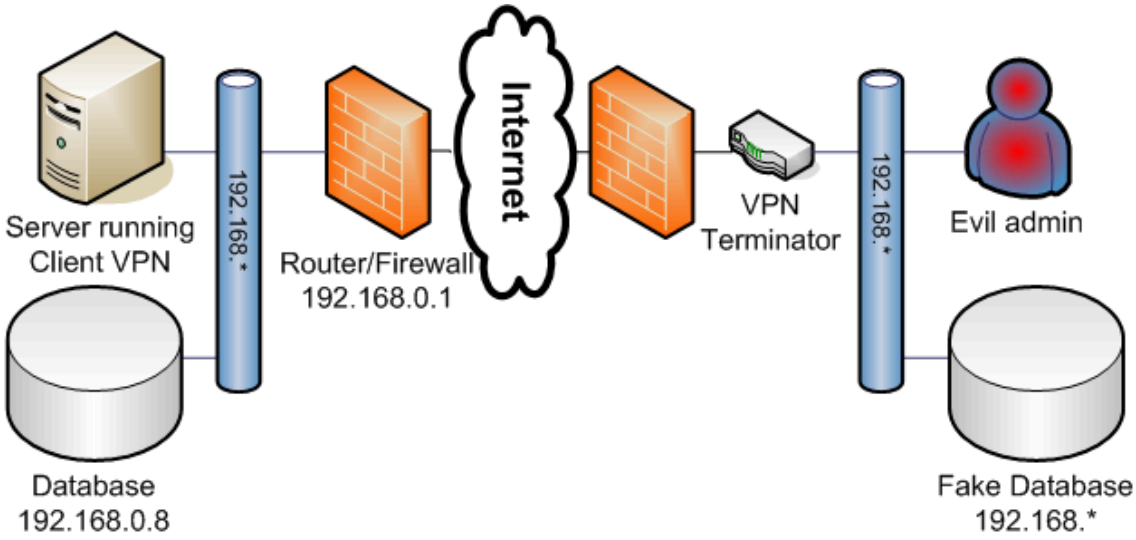

Step 1) Server connects via client VPN to hostile vendor network
Step 2) Evil admin pushes all RFC1918 routes to connect to his own network
Step 3) Server connects to 192.168.0.8 to instantiate a database connection
Step 4) Server connects to evil network instead of the real database (also giving attacker access to the DB password) and potentially sending it sensitive information upon connecting (whatever it would normally request or send upon connecting to the database). Exploits are possible here as well...
Step 5) Evil admin turns off the VPN client forcing the routes to change back to their original settings (depending on the client).
Step 6) Server connects to internal addresses within its own network and re-establishes DB connection
Step 7) Upon successful exploitation Evil Admin allows the server to reconnect to the VPN, but pushes only necessary routes.

Fig 3. Click to enlarge

The third and final VPN issue found in *Fig 3* is between two sets of servers that are interconnected by way of a client VPN. Like the previous examples, the evil administrator can push routes for RFC1918 address space, and cause the remote server to re-route it's traffic over the Internet. This could affect database connections, APIs, email, SMB backups and so on, allowing a remote administrator to temporarily interrupt services, compromise servers and so on.

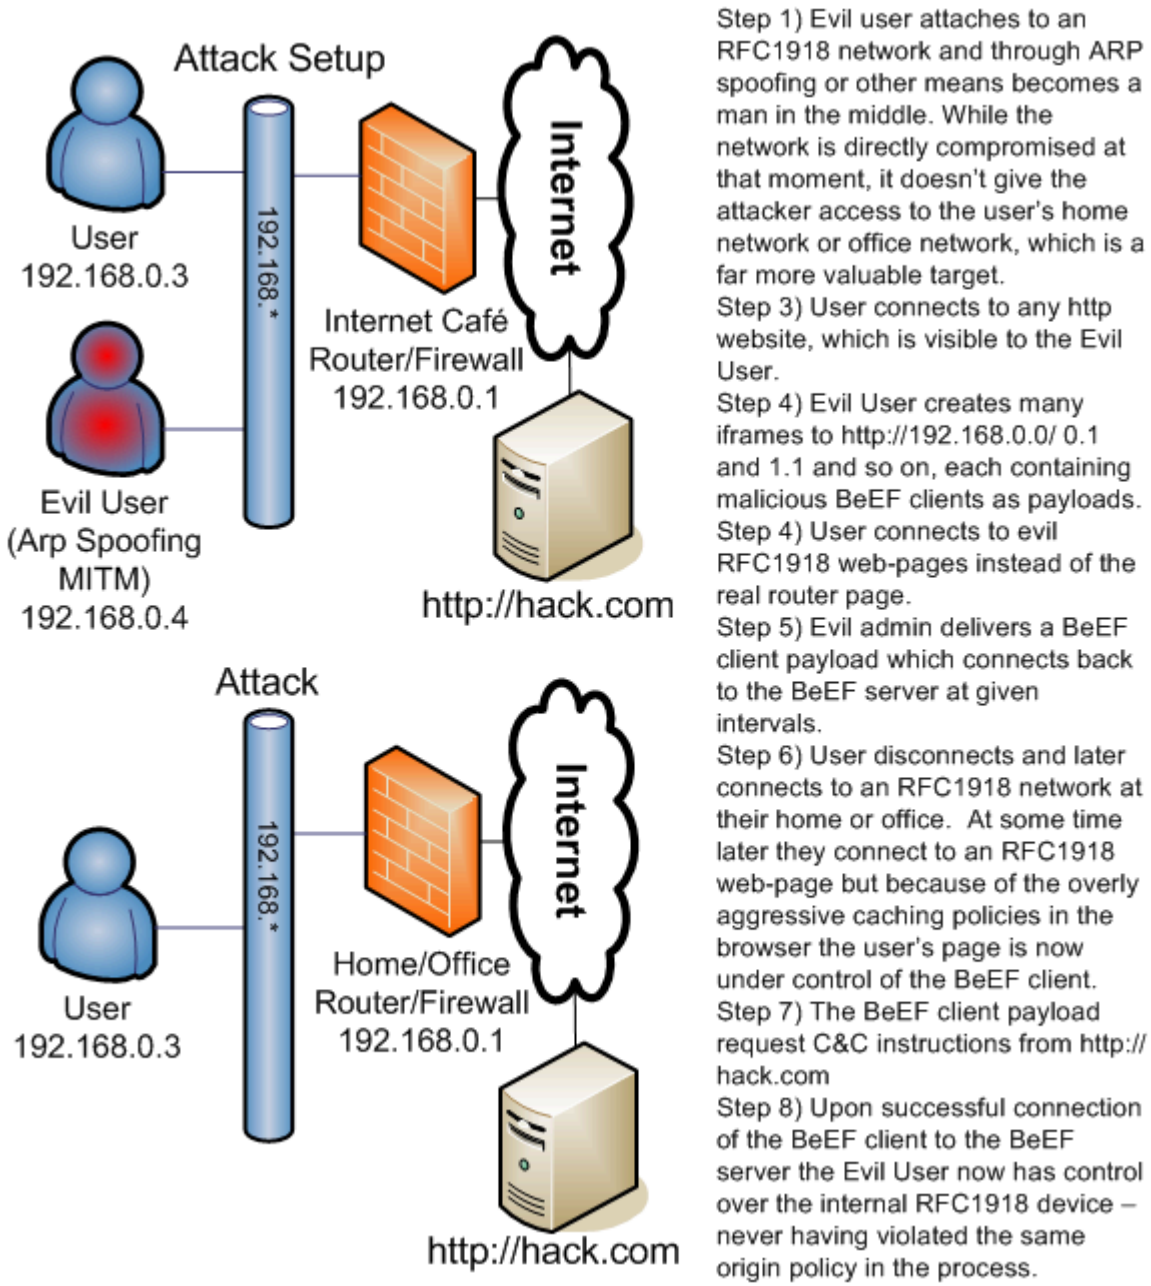

Fig 4. Click to enlarge

Another issue that falls outside of the client VPN issues described above would be a man in the middle attack scenario as seen in *Fig 4*. Most security experts would say once a man in the middle attack is in progress there is little point discussing the issue further, because the user is already completely compromised. While this is somewhat true, it doesn't necessarily give the attacker what they are interested in. For instance an internet cafe may provide the attacker with access to webmail or social networking accounts, but it may not give the attacker access to the user's home network or work network.

An attacker in this scenario could interrupt and modify HTTP requests and inject malicious iframes to the possibly intended RFC1918 targets. Because of the aggressive caching policy within the browser, the malicious JavaScript can be cached well beyond that session. Once the user leaves the Internet Cafe and then turns on their portable device within the context of another RFC1918 network, it is simply a matter of time before the user visits one of these pages, if they are an administrator or someone who has access to physical devices.

In all of these attacks an attacker could do research on all modern networking devices that use JavaScript includes, stylesheets or other objects that can embed JavaScript within them. By forcing these included files to be cached and include the malicious JavaScript client within them, security devices can be thwarted, without necessarily having to know which one the user uses ahead of time, and without having to have them logged in ahead of time (as would be the case with CSRF and DNS rebinding attacks). Instead, the attacker can simply wait, theoretically indefinitely for the attack to work. Realistically the exposure to attack is a limit that is dependent on the duration of the compromised user's cache.

Caveats and Defenses: This attack relies on a number of factors. Firstly, the first three attacks rely on the fact that VPNs can be told what to route. If the VPN can be limited to only route the IP spaces that both parties agree upon, this attack would quickly fall down, or at minimum would only be affective against the IP addresses that were allowed to be routed. All of these attacks require that the browser caches content and that that content persists beyond the initial request.

Additionally, most of these attacks could be thwarted by simply not using actual IP addresses, but rather fully qualified but internal domain names because this would require an attacker to have prior knowledge about the IP to DNS mapping. Also, the use of SSL/TLS on all internal devices would cause a mis-match error if the attacker attempted to cache the JavaScript over HTTPS. Removing all scripting and dynamic content from the browser is also an option although severely limiting as well. Ultimately, most of these issues, aside from the ones found in *Fig 3*, could be mitigated by simply removing persistent cache regularly, or upon the change of any routing information at the operating system level.

Conclusions: Relying entirely upon RFC1918 and built in security functions within the browser to protect users is futile. The browser's same origin policy does not apply if the IP address is the same, and RFC1918 by definition must be the same. VPNs should have an option to define static routes as to mitigate arbitrarily changed routes on the client in a possibly adversarial situation. Ultimately RFC1918 is a poor alternative to IPv6.

Thanks: Special thanks to James Flom for technical oversight and editing help, and to HD Moore and Amit Klein for inspiration and helping me think through some of these ideas.
